

A FULLY DISTRIBUTED ALGORITHM FOR THROUGHPUT PERFORMANCE IN WIRELESS NETWORKS

EYJÓLFUR I. ÁSGEIRSSON, MAGNÚS M. HALLDÓRSSON, AND PRADIPTA MITRA

ABSTRACT. We study link scheduling in wireless networks under stochastic arrival processes of packets, and give an algorithm that achieves stability in the physical (SINR) interference model. The efficiency of such an algorithm is the fraction of the maximum feasible traffic that the algorithm can handle without queues growing indefinitely. Our algorithm achieves two important goals: (i) efficiency is independent of the size of the network, and (ii) the algorithm is fully distributed, i.e., individual nodes need no information about the overall network topology, not even local information.

1. INTRODUCTION

Designing high-performance scheduling algorithms for wireless networks has become an increasingly important topic in recent years. Scheduling in a wireless environment is a non-trivial problem, since simultaneous transmissions interfere with each other in complex ways. A two-fold challenge of appropriately modelling the interference, and then developing algorithms for that model presents itself. Furthermore, in many realistic settings, a centralized controller cannot be assumed, and algorithms that work in a distributed fashion have to be developed.

In this work, we are interested in stability and the associated throughput performance of scheduling algorithms for wireless networks under realistic interference models. We assume that packets arrive at potential senders according to a stochastic process, and the goal of an algorithm is to schedule these transmissions so that the queues of unscheduled packets at each sender remain bounded (in which case, the system is called *stable*). A rich body of research has been devoted to dealing with this issue in a variety of settings. The seminal work of Tassiulas and Ephremides [23] established that an optimal scheduling policy exists, one that stabilizes the system under all arrival rates for which stability is potentially possible. In most settings, however, such a “perfect” solution is computationally intractable, and additionally a distributed implementation is unlikely.

Hence, the search for efficient and/or distributed algorithms which, if not as good as the optimal algorithm, are nevertheless useful. Since these algorithms may not stabilize all feasible arrival processes, the concept of *efficiency* of an algorithm has been introduced, being the fraction of Λ that the algorithm can stabilize, where Λ is the space of arrival processes that the optimum algorithm can stabilize. There have been many approaches to developing such algorithms. A natural step in the search for efficient algorithms is to seek maximal solutions. In the context of wireless networks this is known as Greedy Maximal Scheduling (GMS) algorithm [10] or Longest Queue First (LQF) algorithm [4]. The stability and efficiency of LQF has been investigated extensively [10, 16, 4]. Many other approaches have been proposed as well [18, 11, 22].

Most analytic work on wireless networks has been done in graph-based interference models (e.g. [18, 10, 16, 4, 22]). In these models, wireless links (a link is a sender-receiver pair) that are neighbors in a specified link-graph cannot transmit simultaneously. Though interesting in their own right, these models are known to over-simplify interference coupling [17, 20]. As a result, many research communities working on wireless networks have increased their focus on the so-called *physical model* or the SINR model. In this model, a transmission is considered successful if the signal received at the intended receiver is suitably larger than the cumulative interference due to all other transmissions

in the network, plus the ambient noise. In the SINR model, solving the characterization in [23] is equivalent to solving the maximum weighted *capacity* problem, which is known to be NP-hard [6] and additionally has no known constant factor approximation algorithm (quite apart from the issue of distributed implementation).

In this paper, we develop an algorithm that is completely distributed, with nodes requiring no topological information about the network (not even information about “neighbors”), and achieves an efficiency ratio that is independent of the network size (i.e., the number of links). Thus, the algorithm is scalable in relation to network size. It can also operate in an asynchronous setting, with nodes appearing arbitrarily (as long as the stochastic process meets the required condition).

We give simulation results which lend credence to the theoretical bounds. Our algorithm is extremely robust under fairly high amount of load (achieving efficiency ratios bordering on 0.5).

The only other work (that we are aware of) on stability of algorithms in the SINR model is the recent work by Le *et. al.* [14], who analyze the stability properties of the LQF algorithm. The authors show that the basic LQF is not efficient, but a variation of it that localizes interference is shown to work, with efficiency similar to ours. In a related work, [15] also consider the SINR model, but the model there is different (links are always feasible, but have different data rates based on the SINR achieved. This sort of problem is rather different from the “combinatorial” situation at hand).

A distinguishing feature of our distributed implementations is that they require almost no additional “infrastructure”. Often distributed algorithms for wireless networks have to assume another underlying information infrastructure that can be used to run a localized and/or distributed algorithm, and that infrastructure, moreover, is not subject to the interference constraints of the original network. This is the case with [14], as well as many other works on the topic ([18, 15] for example). This is a rather strong assumption, especially in light of the fact that in a wireless network, one is usually trying to establish such an infrastructure in the first place. It is interesting that we can do without them while obtaining high throughput performance.

From a technical perspective, we adopt the vocabulary and techniques developed in the context of worst-case algorithmic research on the SINR model ([19, 9, 12, 8]). The concept of “affectance” (defined later) developed in some of these works turns out to be quite effective in this context. This approach may have further applications in the study of stability of wireless networks.

The paper is organized as follows. In Section 2, we describe our algorithm and state the main result. In Section 3 we present the system model, and discuss related work further in Sec. 4. The proof of the stability result is given in Section 5. Finally, in Section 6 we present simulation results.

2. ALGORITHM AND RESULT

The wireless network is modeled as a set L of n links, where each link $l_u \in L$ represents a potential transmission from a sender s_u to a receiver r_u , each a point in a metric space.

We assume that packets arrive at the sender of each link l_u according to a stochastic process with average arrival rate m_u .

The extremely simple and fully distributed algorithm is as follows.

Our main result is:

Theorem 1. *For all given networks with links on metric spaces, and all sublinear, length-monotone power assignments, **Reflect** achieves an efficiency ratio independent of n .*

3. SOME PRELIMINARIES

The distance between two points x and y is denoted $d(x, y)$. The distance from l_u ’s sender to l_v ’s receiver is denoted $d_{uv} = d(s_u, r_v)$. The length of link l_u is denoted simply by $\ell_u = d(s_u, r_u)$. The link set is associated with a *power assignment* P , which is an assignment of a transmission power

Algorithm 1 Reflect (Run by each link l_u in the system)

```

1:  $Q_u \leftarrow \emptyset$  (queue of outstanding packets)
2: for  $t \leftarrow 1, 2, \dots$  do
3:   Let  $\mathcal{A}$  be the set of packets that arrive at the beginning of time slot  $t$ 
4:   Add  $\mathcal{A}$  to the end of  $Q_u$ 
5:   if  $Q_u$  is non-empty then
6:     Transmit a packet from  $Q_u$  with probability  $2.5 \cdot m_u$ 
7:   end if
8: end for

```

P_u to be used by the sender of each link $l_u \in L$. The signal received at point y from a sender at point x with power P_x is $P_x/d(x, y)^\alpha$ where the constant $\alpha > 0$ is the *path-loss exponent*.

We can now describe the *physical* or SINR-model of interference. In this model, a receiver r_u successfully receives a message from the sender s_u if and only if the following condition holds:

$$(1) \quad \frac{P_u/\ell_u^\alpha}{\sum_{l_v \in S \setminus \{l_u\}} P_v/d_{vu}^\alpha + N} \geq \beta ,$$

where N is the environmental noise, the constant β denotes the minimum SINR (signal-to-interference-noise-ratio) required for a message to be successfully received, and S is the set of concurrently scheduled links in the same *slot* (we assume that time is slotted). We say that S is *SINR-feasible* (or simply *feasible*) if (1) is satisfied for each link in S .

A power assignment P is *length-monotone* if $P_v \geq P_w$ whenever $\ell_v \geq \ell_w$ and *sublinear* if $\frac{P_v}{\ell_v^\alpha} \leq \frac{P_w}{\ell_w^\alpha}$ whenever $\ell_v \geq \ell_w$. This class includes the most interesting and practical power assignments, such as uniform power (all links use the same power), linear power ($P_u = \ell_u^\alpha$, known to be energy efficient in the presence of noise), and mean power ($P_u = \ell_u^{\alpha/2}$, the assignment that produces maximum capacity in this class [8]). Let $\Delta = \frac{\ell_{\max}}{\ell_{\min}}$ where ℓ_{\max} and ℓ_{\min} are, respectively, the maximum and minimum lengths of links in L .

Definition 2. The affectance $a_v^P(u)$ of link l_u caused by another link l_v , with a given power assignment P , is the interference of l_v on l_u relative to the signal received, or

$$a_v^P(u) = \min \left\{ 1, c_u \frac{P_v}{P_u} \cdot \left(\frac{\ell_u}{d_{vu}} \right)^\alpha \right\} ,$$

where $c_u = \beta/(1 - \beta N \ell_u^\alpha / P_u)$.

We need the following assumption.

Assumption 3. $c_u \leq 2\beta$ for any link l_u .

This is a fairly reasonable assumption. It simply says that in the absence of other links, the transmission succeeds comfortably. The constant 2 is not fundamental; any value greater than 1 would suffice.

Since $c_u \geq \beta$ by definition, this implies that

$$(2) \quad \frac{c_v}{c_u} \leq 2 \text{ for any two links } l_u, l_v$$

The definition of affectance was introduced in [5, 9] and achieved the form we use in [13]. When clear from the context we drop the superscript P . Also, let $a_v^P(v) = 0$. Using affectance, Eqn. 1 can be rewritten as

$$(3) \quad a_S^P(u) \equiv \sum_{l_v \in S} a_v^P(u) \leq 1 ,$$

for all $l_u \in S$.

Signal-strength and robustness. A δ -signal set of links is a set of links where the affectance on any link is at most $1/\delta$. A set is SINR-feasible iff it is a 1-signal set. We know:

Lemma 4 ([7]). *Let ℓ_u, ℓ_v be links in a q^α -signal set. Then, $d_{uv} \cdot d_{vu} \geq q^2 \cdot \ell_u \ell_v$.*

Now we switch to the aspects of this paper related to queueing theory and stochastic processes. We first define stability.

Definition 5. *An algorithm stabilizes a network for a particular arrival process if, under that arrival process, the average queue size is bounded (at any given time).*

The *throughput region* is then the set of all possible arrival rate vectors such that there exists some scheduling policy that can stabilize the network.

As proved in [23], the *throughput region* is characterized by

$$\Lambda = \{\lambda : \lambda \preceq \phi, \text{ for some } \phi \in Co(\Omega)\} ,$$

where Ω is the set of all maximal feasible schedules (meaning arrival processes that put weight 1 on a single maximal feasible set) and $Co(\Omega)$ is the convex hull of Ω . Note that λ and ϕ are n -dimensional vectors and $\lambda \preceq \phi$ means each element of λ is upper bounded by the corresponding one in ϕ .

Since fast and/or distributed algorithms might not stabilize all of Λ , one hopes to achieve a large *efficiency ratio*.

Definition 6. *The efficiency ratio γ of a scheduling algorithm is $\gamma = \sup\{\eta : \text{all networks are stabilized for all } \lambda \in \eta\Lambda\}$, where $\eta\Lambda = \{\eta\lambda : \lambda \in \Lambda\}$.*

We assume that the arrival process on a link is i.i.d. across time, and different links are independent of each other.

We will use M_i to denote both maximal feasible sets, and characteristic vectors of said sets (the usage being clear from context). For a given efficiency ratio γ , it must hold for all permissible arrival rate vectors λ that $\lambda \preceq \sum_i m_i M_i$, where m_i are weights such that

$$(4) \quad \sum_i m_i = \gamma .$$

It can be easily seen that for any link l_u ,

$$(5) \quad m_u = \sum_{i: l_u \in M_i} m_i \leq \gamma .$$

4. RELATED WORK

As stated in Thm. 1, the efficiency of the algorithm is independent of n (the number of links in the system). It is, however, dependent on another network parameter Δ , the ratio between the longest and the shortest link in the system (the proof in the next section contains the exact expression). The only comparable work on this model [14] has the same dependence on Δ (this is not explicitly stated in the paper, but can be seen to be necessary). The main discriminating feature of our work is that it is distributed in a much stronger sense. The algorithm in [14] can be characterized as “localized”, where each link needs to be aware of and have communicated with other links in its neighborhood. We have no need for such infrastructure.

In terms of efficiency ratio, a range of results have been derived in a variety of models. Naturally one seeks efficiency of 1 whenever possible [18], but results for efficiency ratio of 1 under certain conditions [4], or $\frac{1}{2}$ [3], or $\frac{1}{6}$ [10] can be found in the literature. Ratios in terms of certain network characteristics are known as well – such as in terms of the degree of the interference graph [2] or

the local pooling factor [10, 14]. In [21], the abstract SINR model (received signal is a general function instead of being length-based) in the context of MIMO networks is studied. An efficiency ratio based on a system-specific value (“effective interference number”) is derived, with no direct comparison with distance-based SINR models.

For the SINR model, an efficiency ratio that is an “unconditional” constant (independent of both n and Δ) is not known.

5. PROOF OF STABILITY

We now present a proof of Thm. 1.

Note that the probability $2.5 \cdot m_u$ used for link l_u in the algorithm is well-defined, since we claim stability with constant efficiency ratio bounded from above by $\frac{1}{3}$.

We first need the following observation.

Observation 7. *For any two links ℓ_u and ℓ_v using a length monotone, sub-linear power assignment,*

$$\frac{P_u \ell_v^\alpha}{P_v \ell_u^\alpha} \leq \Delta^\alpha .$$

Proof. If $\ell_v \leq \ell_u$, it holds by sub-linearity that $P_u/\ell_u^\alpha \leq P_v/\ell_v^\alpha$. Otherwise, if $\ell_u \leq \ell_v$, then by monotonicity $P_u \leq P_v$ and by definition of Δ , $\ell_v^\alpha/\ell_u^\alpha \leq \Delta^\alpha$. \square

The following key lemma shows that no link is affected too much by any single feasible set.

Lemma 8. *Consider a feasible set S and a link l_v (not necessarily a member of S). Then,*

$$(6) \quad \sum_{l_z \in S} a_z(v) \leq \kappa \cdot \Delta^\alpha ,$$

for some constant κ .

Proof. We use the signal strengthening technique of [9]. For this, we decompose the set S to $\lceil 2 \cdot 3^\alpha / \beta \rceil^2$ sets, each a 3^α -signal set. We prove the claim for one such set; since there are only constantly many such sets, the overall claim holds (with the appropriate increase in the constant factor). Let us reuse the notation S to be such a 3^α -signal set.

Consider the link $l_u = (s_u, r_u) \in S$ such that $d(r_v, r_u)$ is minimum. Also consider the link $l_w = (s_w, r_w) \in S$ such that $d(s_w, r_v)$ is minimum. Let $D = d(r_v, r_u)$. We claim that for all links $l_x = (s_x, r_x) \in S$ with $\ell_x \neq \ell_w$, it holds that

$$(7) \quad d(s_x, r_v) \geq \frac{1}{2} D .$$

To prove this, assume, for contradiction, that $d(s_x, r_v) < \frac{1}{2} D$. Then, $d(s_w, r_v) < \frac{1}{2} D$, by definition of l_w . Now, again by the definition of l_u , $d(r_x, r_v) \geq D$ and $d(r_w, r_v) \geq D$. Thus $\ell_w \geq d(r_w, r_v) - d(r_v, s_w) > \frac{D}{2}$ and similarly $\ell_x > \frac{D}{2}$. On the other hand $d(s_w, s_x) < \frac{D}{2} + \frac{D}{2} = D$. Now, $d_{wx} \cdot d_{xw} \leq (\ell_w + d(s_w, s_x))(\ell_x + d(s_w, s_x)) < (\ell_w + D)(\ell_x + D) < 9\ell_w \ell_x$, contradicting Lemma 4.

Now that we have proven Eqn 7, by the triangle inequality, $d_{xu} = d(s_x, r_u) \leq d(s_x, r_v) + d(r_v, r_u) \leq 3d(r_v, s_x) = 3d_{xv}$. Applying Obs. 7, we see that

$$\frac{a_x(v)}{a_x(u)} \leq \frac{c_v}{c_u} \cdot \frac{P_u}{P_v} \cdot \frac{d_{xu}^\alpha}{d_{xv}^\alpha} \cdot \frac{\ell_v^\alpha}{\ell_u^\alpha} \leq 2 \cdot 3^\alpha \frac{P_u \ell_v^\alpha}{P_v \ell_u^\alpha} \leq 2 \cdot 3^\alpha \Delta^\alpha .$$

where $\frac{c_v}{c_u} \leq 2$ follows from Eqn. 2. Finally, summing over all links in S ,

$$\begin{aligned} a_S(v) &= \sum_{l_z \in S} a_z(v) = a_w(v) + \sum_{l_z \in S \setminus \{l_w\}} a_z(v) \\ &\leq 1 + 2 \cdot (3\Delta)^\alpha \sum_{l_z \in S \setminus \{l_w\}} a_z(u) \\ &\leq 1 + 2 \cdot (3\Delta)^\alpha, \end{aligned}$$

where we use $a_w(v) \leq 1$ by the definition of affectance, and $\sum_{l_z \in S \setminus \{l_w\}} a_z(u) \leq 1$ since S is feasible and $l_u \in S$.

This completes the proof setting $\kappa = 3^{\alpha+1}$. \square

We turn to the proof of Thm. 1.

Proof. We claim an efficiency ratio of $1/(6\kappa\Delta^\alpha)$ where κ is the constant from Lemma 8. Thus, it is enough to prove stability for all stochastic processes for which the following holds:

$$(8) \quad \sum_{i: M_i} m_i \leq \frac{1}{6\kappa\Delta^\alpha}.$$

Consider the affectance on any link l_u during the execution of the algorithm in a single slot, and denote it by $a(u)$. This can be computed as $a(u) = \sum_{l_v} X_v a_v(u)$, where X_v is a Bernoulli random variable which is 1 iff link l_v has a non-empty queue and chooses to transmit during the same slot. Now,

$$\begin{aligned} \mathbb{E}(a(u)) &= \sum_{l_v} \mathbb{E}(X_v) a_v(u) \stackrel{1}{\leq} \sum_{l_v} 2.5 m_v a_v(u) \\ &\stackrel{2}{=} 2.5 \sum_i m_i \sum_{M_i \ni l_v} a_v(u) \\ &\stackrel{3}{\leq} 2.5 \sum_i m_i \kappa \Delta^\alpha \stackrel{4}{\leq} \frac{2.5 \kappa \Delta^\alpha}{6\kappa \Delta^\alpha} = \frac{5}{12}, \end{aligned}$$

where explanations of the numbered (in)equalities are:

- (1) $\mathbb{E}(X_v) \leq 2.5 m_v$ by the description of the algorithm.
- (2) By Eqn. 5 and rearrangement.
- (3) By Lemma 8.
- (4) By Eqn. 8.

Thus, with probability at least $\frac{1}{2}$, $a(u) < 1$ (by Markov's inequality). Hence, if l_u has a non-empty queue, with probability at least $\frac{5}{4} \cdot m_u$, the queue size decreases. Note that the probability can potentially be higher than $\frac{5}{4} \cdot m_u$, but never smaller. Therefore, this system is at least as efficient as the system where in each slot the queue size reduces by 1 with i.i.d. probability exactly $\frac{5}{4} \cdot m_u$.

The queue dynamics on a single link become equivalent to the following single server system with slotted time and an infinite queue. In this system, at the beginning of each time slot, A packets arrive, where A is a random variable on the non-negative integers, with $\mathbb{E}(A) = m_u$. At the end of each slot, the server processes D packets (or empties the queue), where D is a Bernoulli random variable with $\mathbb{E}(D) = \frac{5}{4} \cdot m_u > m_u = \mathbb{E}(A)$. Since the departure process is faster than the arrival process, the stability of the queue is guaranteed by basic results in queueing theory [1]. \square

We note that for linear power assignment, the dependence on Δ in Obs. 7 completely disappears, and thus also in Lemma 8 and Thm. 1.




FIGURE 1. The maximum queue lengths for the distributed algorithms Distr-SingleLink and Reflect. The problem instances are based on random topology with $n = 200$, $\ell_{\min} = 1$, $\ell_{\max} = 20$, $\alpha = 2.5$ and $\beta = 1$.

Corollary 9. *For all given networks with links on metric spaces, using linear power assignment, **Reflect** achieves an efficiency ratio that is an absolute constant (independent of n and Δ).*

Remark: In **Reflect** we have assumed that each link l_u knows m_u . In practice, this can be easily approximated at time t by $\min\{1, \frac{A(t)}{t}\}$ (where $A(t)$ is the number of packet arrivals up to time t), which converges to the right value almost surely.

5.1. Link partitioning. If link lengths are known beforehand and some pre-processing is allowed, the efficiency can be made to have a better dependence on Δ , specifically, we can achieve an efficiency ratio of $1/(6 \cdot 2^\alpha \kappa \log \Delta)$.

We can partition the link set into a collection of nearly equi-length link sets, i.e., sets where the lengths in the set vary by at most a factor of 2. It is easy to show that a link set L can be partitioned into at most $\log_2 \Delta + 1$ sets of nearly equi-length links L_r for $r = 1 \dots \log_2 \Delta + 1$ where L_r contains links of lengths in $[2^{r-1} \cdot \ell_{\min}, 2^r \cdot \ell_{\min})$.

We partition the time slots accordingly, a time slot t is used to schedule links from class L_r where $r = ((t - 1) \bmod (\log_2 \Delta + 1)) + 1$.

With the partition, the arrival process on a sequence of slots devoted to a single length class is equivalent to the setting where all links are nearly equi-length and assuming that $\sum_i m_i = 1/(6 \cdot 2^\alpha \kappa)$.¹ This partitioning combined with Lemma 8 proves the claimed efficiency.

6. SIMULATIONS

To see how the distributed algorithm **Reflect** performs, we ran simulations on instances based on random topology. The problem instances were created by generating random links in a rectangle with side length 100. The length of the links were uniform random variables between ℓ_{\min} and ℓ_{\max} , which we set as 1 and 20 respectively. We generated random transmission requests for 100,000 time slots while running the algorithms. Our focus was on the behavior of the maximum queue length, i.e., the largest number of waiting transmission requests over all the links, which was measured every 10,000 rounds.

Figure 1 shows the results of the distributed algorithm **Reflect** for random instances with 200 links after 10,000, 50,000 and 100,000 time slots. The probability scaling factor, ρ , defines the load

¹There are certain technicalities here, since a) $M_i \cap L_t$ may not be a *maximal* feasible set in L_t and b) Two sets M_i and M_j may have the same “projection” in L_t , i.e., it is possible that $M_i \cap L_t = M_j \cap L_t$ for $M_i \neq M_j$. These can be handled in a straightforward way.

on the system. Thus when $\rho = 1$, the system will, in expectation, receive a maximal feasible set in every round, while if the scaling factor is zero, no requests are generated. The efficiency ratio of an algorithm is then equal to the largest ρ such that the algorithm is stable. We used a granularity of 0.01 for values of ρ between 0.01 and 0.6 and took the average over 10 runs for each value of ρ .

The **Reflect** algorithm used in Figure 1 approximates the arrival rate of requests for each link, as mentioned in an earlier remark, instead of assuming that the links know the arrival rate of requests.

It is interesting to note, in Figure 1, that there is a very sharp threshold where the algorithm is no longer stable. As soon as the algorithm becomes unstable, the queue lengths increase very rapidly. As expected, the centralized Longest Queue First (LQF) algorithm was more stable than the distributed algorithm, managing to keep the maximum queue length below 2 for all values of $\rho \leq 0.6$ and not becoming unstable until $\rho > 0.9$. We note here again that though our algorithm has a lower (though still high) throughput, its main feature is the lack of a requirement for a centralized or even localized control.

REFERENCES

- [1] Søren Asmussen. *Applied Probability and Queues*. Springer, 2nd edition, 2003.
- [2] Prasanna Chaporkar, Koushik Kar, Xiang Luo, and Saswati Sarkar. Throughput and fairness guarantees through maximal scheduling in wireless networks. *IEEE Transactions on Information Theory*, 54(2):572–594, 2008.
- [3] J. G. Dai and Balaji Prabhakar. The throughput of data switches with and without speedup. In *INFOCOM*, pages 556–564, 2000.
- [4] Antonis Dimakis and Jean Walrand. Sufficient conditions for stability of longest-queue-first scheduling: second-order properties using fluid limits. *Advances in Applied Probability*, 38(2):505–521, 2006.
- [5] Olga Goussevskaia, Magnús M. Halldórsson, Roger Wattenhofer, and Emo Welzl. Capacity of Arbitrary Wireless Networks. In *INFOCOM*, pages 1872–1880, April 2009.
- [6] Olga Goussevskaia, Yvonne A. Oswald, and Roger Wattenhofer. Complexity in Geometric SINR. In *Mobihoc*, pages 100–109, 2007.
- [7] M. M. Halldórsson. Wireless scheduling with power control. *ACM Transactions on Algorithms*. To appear. See also <http://arxiv.org/abs/1010.3427>, September 2010.
- [8] Magnús M. Halldórsson and Pradipta Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *SODA*, 2011.
- [9] Magnús M. Halldórsson and Roger Wattenhofer. Wireless Communication is in APX. In *ICALP*, pages 525–536, July 2009.
- [10] Changhee Joo, Xiaojun Lin, and N.B. Shroff. Understanding the Capacity Region of the Greedy Maximal Scheduling Algorithm in Multi-Hop Wireless Networks. In *INFOCOM*, 2008.
- [11] Changhee Joo and Ness B. Shroff. Performance of random access scheduling schemes in multi-hop wireless networks. *IEEE/ACM Trans. Netw.*, 17(5):1481–1493, 2009.
- [12] Thomas Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *SODA*, 2011.
- [13] Thomas Kesselheim and Berthold Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178, August 2010.
- [14] Long B. Le, Eytan Modiano, Changhee Joo, and Ness B. Shroff. Longest-queue-first scheduling under SINR interference model. In *MOBIHOC*, 2010.
- [15] Hyang-Won Lee, Eytan Modiano, and Long Bao Le. Distributed throughput maximization in wireless networks via random power allocation. In *IEEE WiOpt*, August 2009.
- [16] Bo Li, Cem Boyaci, and Ye Xia. A refined performance characterization of longest-queue-first policy in wireless networks. In *MobiHoc*, pages 65–74, 2009.
- [17] Ritesh Maheshwari, Shweta Jain, and Samir R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys*, pages 141–154, 2008.
- [18] Eytan Modiano, Devavrat Shah, and Gil Zussman. Maximizing throughput in wireless networks via gossiping. In Raymond A. Marie, Peter B. Key, and Evgenia Smirni, editors, *SIGMETRICS/Performance*, pages 27–38. ACM, 2006.
- [19] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *INFOCOM*, 2006.
- [20] Thomas Moscibroda, Roger Wattenhofer, and Yves Weber. Protocol Design Beyond Graph-Based Models. In *Hotnets*, November 2006.
- [21] D. Qian, D. Zheng, J. Zhang, and Shroff N. CSMA-based distributed scheduling in multi-hop MIMO networks under SINR model. In *IEEE INFOCOM*, 2010.
- [22] Devavrat Shah, Jinwoo Shin, and Prasad Tetali. Efficient distributed medium access (to appear). In *FOCS*, 2011.
- [23] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.*, 37(12):1936–1948, 1992.

(M. M. Halldórsson) SCHOOL OF SCIENCE AND ENGINEERING, REYKJAVIK UNIVERSITY, 101 REYKJAVIK, ICELAND,
E-mail address: eyjo@ru.is

(M. M. Halldórsson) SCHOOL OF COMPUTER SCIENCE, REYKJAVIK UNIVERSITY, 101 REYKJAVIK, ICELAND,
E-mail address: mmh@ru.is

(P. Mitra) SCHOOL OF COMPUTER SCIENCE, REYKJAVIK UNIVERSITY, REYKJAVIK 101, ICELAND
E-mail address: ppmitra@gmail.com